

The UltraSPARC T1 Processor - High Bandwidth For Throughput Computing

William Bryg, Distinguished Engineer
Jerome Alabado, Staff Engineer
Sun Microsystems, Inc.

December 2005



Sun Microsystems, Inc.
4150 Network Circle
Santa Clara, CA 95045 U.S.A.
www.sun.com

Copyright 2005 Sun Microsystems, Inc. 4150 Network Circle, Santa Clara, California 95045 U.S.A. All rights reserved.

Sun Microsystems, Inc. has intellectual property rights relating to technology described in this document. In particular, and without limitation, these intellectual property rights may include one or more patents or pending patent applications in the U.S. or other countries.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the United States and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, Java, and Solaris are trademarks or registered trademarks of Sun Microsystems, Inc. in the United States and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the US and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

U.S. Government Rights—Commercial use. Government users are subject to the Sun Microsystems, Inc. standard license agreement and applicable provisions of the FAR and its supplements.

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2005 Sun Microsystems, Inc., 4150 Network Circle, Santa Clara, Californie 95045 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque enregistrée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company Ltd.

Sun, Sun Microsystems, le logo Sun, Java, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REPOUDRE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.

1.0 Introduction

Companies today face significant computing challenges. The demand for network services continues to rise while IT budgets and resources remain limited. Sun's response to this growing problem is innovation. Its Chip Multithreading (CMT) processor strategy for throughput computing will usher in a computing revolution that will dramatically reduce the cost and complexity of network computing by delivering performance and power efficiency that are orders of magnitude greater than contemporary processors.

An integral part of a balanced systems approach is providing sufficient bandwidth.

The first instance of a radical CMT processor is the UltraSPARC T1 processor. The processor's power efficiency and performance advantages are a result of Sun's long history and design philosophy of building well-balanced systems. An integral part of this balanced approach is an emphasis that a system must have sufficient bandwidth. CMT processors are capable of executing many instructions simultaneously, so the infrastructure inside and around the processor must ensure that data is available when needed. With this notion in mind, Sun designed the UltraSPARC T1 processor capable of sustaining high bandwidths.

2.0 UltraSPARC T1 Processor: A High Bandwidth Processor

Commercial server workloads have poor data locality and significant data sharing, which can lead to high cache miss rates. When the processor does not find a data item it needs in the cache, a request is propagated to the memory system to fill the missing cache line. The time required to retrieve the requested data depends on both the time it takes to reach memory (latency) and the amount of data moved per unit time (bandwidth). As processor and memory performance have historically diverged, however, retrieving data from memory has become a major bottleneck. To address this problem, Sun designed the UltraSPARC T1 processor to execute many threads in parallel that effectively hides memory latency while fully exploiting a high bandwidth memory system,

Memory is a major bottleneck in commercial server workloads. To address this problem, Sun designed the UltraSPARC T1 processor with the proper balance of many threads and sufficient bandwidth.

ensuring the proper balance between many threads and sufficient bandwidth.

The UltraSPARC T1 processor can retire up to 8 instructions per cycle (IPC). Its crossbar interconnect and L2 cache, along with other processor resources, provide a continuous flow of data into the eight cores.

An example of the UltraSPARC T1 processor's balanced approach is the interaction among its processor cores and associated L1 caches, crossbar interconnect, and L2 cache. Together, along with other processor resources, these enable the UltraSPARC T1 processor to retire up to eight instructions per cycle (IPC), or one instruction per core. For example, if all eight cores are executing instructions and come upon an L1 data cache miss in the same cycle, these eight cache miss requests¹ are sent through the crossbar interconnect switch, which is able to service all eight requests in parallel. The requests are then sent to the L2 cache, where the L2 banks are pipelined and can each service one cache access per bank per cycle. Meanwhile, the availability of many threads prevents the UltraSPARC T1 processor from stalling its pipelines. This seamless exchange between the crossbar interconnect and the L2 cache provide a continuous flow of data into the eight cores.

In addition to the crossbar interconnect and L2 cache, the UltraSPARC T1 processor's memory interface contributes to achieving high IPC.

Another processor resource that contributes to the UltraSPARC T1 processor's high IPC is the memory interface. Like the roles of the crossbar interconnect and the L2 cache in the previous example, the UltraSPARC T1 processor's four on-chip memory controllers ensure data is available when needed by the cores. The UltraSPARC T1 processor, with each of its eight cores supporting four threads, supports a total of 32 threads and effectively hides memory latency by combining ideas from chip multiprocessors and fine-grained multithreading as shown in Figure 1.

1. On average, only two cache miss requests are sent to the interconnect per cycle.

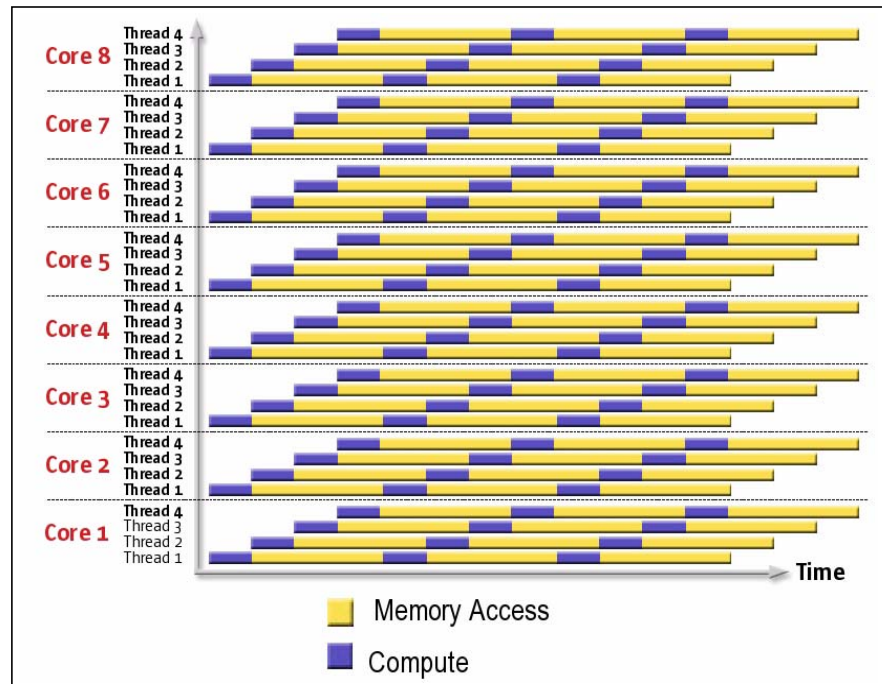


FIGURE 1 The UltraSPARC T1 Processor's Multithreaded Conceptual Diagram

The UltraSPARC T1 processor's memory bandwidth satisfies the required bandwidth demanded by its many threads.

The UltraSPARC T1 processor's ability to execute many threads in parallel, however, places a heavy demand on the memory system. In the case of Figure 1, when a cache miss occurs for thread 1, it is assumed that data is readily available for the next thread (thread 2), the following thread (thread 3), and so on. However, suppose all 32 threads do not have their data readily available. In this worst-case scenario, data must be brought from main memory to the on-chip caches, consuming a total bandwidth of about 20.5 Gbytes/s¹. In most cases, however, all 32 threads will not access main memory every cycle, resulting in a much lower bandwidth consumption. With the UltraSPARC T1 processor's four on-chip memory controllers supporting over 20 Gbytes/s of memory bandwidth, the proper balance between required bandwidth and the bandwidth that the UltraSPARC T1 processor can actually deliver is satisfied.

Several experiments also show that the UltraSPARC T1 processor's memory bandwidth satisfies the bandwidth demanded by many workloads. For example, memory bus statistics taken when running an OLTP (online

1. Total bandwidth consumed = [(cache line size/memory latency) * 32 threads]; [(64 bytes/100 ns) * 32 threads.

transaction processing) test on a the UltraSPARC T1 processor-based system show that its available bandwidth more than satisfies the total bandwidth consumed. In one instance, where the OLTP test completely ran in main memory with minimal I/O accesses, the total bandwidth consumed was 5.4 Gbytes/s, as shown in Table 1. Each on-chip memory controller performed a certain number of reads and writes, for a total of 3.6 Gbytes/s of read bandwidth and 1.9 Gbytes/s of write bandwidth.

TABLE 1 Fully-Cached OLTP Run, Reduced I/O

Fully Cached OLTP Run	Reads/Sec. (in millions)	Writes/Sec. (in millions)	Bandwidth Consumed
DRAM Controller 0	16	8	Reads: 3.6 Gbytes/s Writes: 1.8 Gbytes/s Total: 5.4 Gbytes/s
DRAM Controller 1	14.3	7.2	
DRAM Controller 2	15.4	7.6	
DRAM Controller 3	14.6	7.5	

The UltraSPARC T1 processor’s memory bandwidth more than satisfies the required bandwidth demanded by an OLTP test.

In a more I/O-intensive instance of the test, where the test environment involved more interactions with both memory and disk, total bandwidth consumed increased to 6.15 Gbytes/s, as shown in Table 2.

TABLE 2 Memory and Disk Accesses OLTP Run, 60% of Necessary I/O

Fully Cached OLTP Run	Reads/Sec. (in millions)	Writes/Sec. (in millions)	Bandwidth Consumed
DRAM Controller 0	17	9/6	Reads: 3.95 Gbytes/s Writes: 2.2 Gbytes/s Total: 6.15 Gbytes/s
DRAM Controller 1	16	9.2	
DRAM Controller 2	18	9.3	
DRAM Controller 3	15.3	8.9	

Other enhancements in the UltraSPARC T1 processor enable high IPC, maximize bandwidth, and avoid bottlenecks.

With the UltraSPARC T1 processor supporting over 20 Gbytes/s of memory bandwidth, not only does it satisfy requirements in both test environments, but also leaves plenty of bandwidth headroom for applications to take advantage of and achieve even higher performance.

In addition to the features of the memory interface, L2 cache, and crossbar interconnect that enable the UltraSPARC T1 processor to retire eight instructions per cycle, other features implemented in the processor to maximize bandwidth and avoid bottlenecks include the following:

- **Processor Core and L1 Caches**

The core’s pipeline and L1 caches have been designed from the ground-up to work with other processor resources to maximize bandwidth. Each

core has its own L1 instruction and data caches to provide recently-used data when needed.

- **Crossbar Interconnect Switch**

The crossbar interconnect provides the communication link among the 8 cores, the L2 cache banks, and other shared resources on the processor; running at the core clock-frequency, it provides over 200 Gbytes/s of bandwidth. The crossbar interconnect is also non-blocking, allowing multiple transactions to be queued up from the source, maximizing bandwidth utilization.

- **L2 Cache**

The L2 cache has a directory of all eight L1 caches and has the ability to invalidate cache lines that are modified. This action simplifies coherency checking across the eight cores and conserves bandwidth by eliminating snoop traffic across the chip.

- **Memory Interface**

The four on-chip channels of dual data rate 2 (DDR2) DRAM enable the UltraSPARC T1 processor to directly access memory without having to go through an external chipset. Additionally, with memory closer to the chip, the UltraSPARC T1 processor has less latency to hide than processors with external memory controllers. Each on-chip memory controller also supports outstanding transactions to complete out-of-order with respect to the original request order. This feature enables the memory banks in the system to service requests as soon as a bank is available, maximizing bandwidth utilization. Because most accesses to memory are predominantly reads, the memory interface also gives priority to reads over writes, another feature that maximizes bandwidth utilization.

- **Changes in the I/O Interface**

The I/O interface has been enhanced for more efficient use of JBUS's¹ available bandwidth of 3.2 Gbytes/s. With the interprocessor coherency all done on-chip, JBUS is now reserved for I/O related traffic only, primarily for DMA transactions.

3.0 The Competitive Landscape

No other modern server processor in the industry has gone through a systematic and radical design like the UltraSPARC T1 processor. From a memory bandwidth perspective, the UltraSPARC T1 processor offers over

1. Sun designed the JBUS as an interconnect for systems based on the UltraSPARC III processor. JBUS supports up to 4-processor configurations, delivering up to 3.2 GB/s of bandwidth. In multiprocessor configurations, coherence traffic dominated JBUS cycles. In the case of the UltraSPARC T1 processor, the entire JBUS is dedicated for I/O traffic, maximizing the efficient use of the available bandwidth provided by the JBUS design.

three times the memory bandwidth of contemporary server processors, as listed in Table 3.

TABLE 3 Memory Bandwidth of Contemporary Processors

Processor	Memory Bandwidth (in Gbytes/s)	Comment
Sun UltraSPARC T1 Processor	>20	4 on-chip memory controllers
Intel Xeon	6.4 ^a	Requires an external chipset, shared bus
Intel Itanium 2	6.4 ^b	Requires an external chipset, shared bus
AMD Opteron	6.4 ^c	On-chip memory controller

a. Intel Xeon Datasheet, <ftp://download.intel.com/design/Xeon/datashts/30235501.pdf>; Intel's Xeon memory bandwidth decreases significantly as processors are added to a multiprocessor configuration. 6.4 Gbytes/s memory bandwidth is available only for a single-processor configuration.

b. Intel E8870 Chipset Overview, <http://www.intel.com/design/chipsets/e8870/textonly.htm>

c. AMD Opteron Processor Key Architectural Features, http://www.amd.com/us-en/Processors/ProductInformation/0,,30_118_8796_8805,00.html

The UltraSPARC T1 processor offers over 3 times the memory bandwidth of Intel's Xeon and Itanium processors.

Moreover, the UltraSPARC T1 processor's memory bandwidth in Table 3 reflects overhead that the processor incurs as a result of normal operation. In contrast, memory bandwidth for all other processors is noted as a theoretical maximum. Actual memory bandwidth achieved by these processors will vary depending on the application. In an analysis by a hardware enthusiast¹, the actual Opteron memory bandwidth delivered was 5.7 Gbytes/s out of a theoretical 6.4 Gbytes/s, while the Pentium 4 fared even worse - delivering a 4.4 Gbytes/s memory bandwidth out of a theoretical 6.4 Gbytes/s. Moreover, the Pentium 4 could not even sustain 4 Gbytes/s when the board was set to normal settings.

The UltraSPARC T1 processor's on-chip memory controllers enable memory requests to travel much shorter distances than Intel's Xeon and Itanium processors.

The UltraSPARC T1 processor's on-chip memory controllers also enable memory requests to travel much shorter distances than Intel's Xeon and Itanium processors, which must travel through an external chipset instead of directly accessing main memory. This extra distance not only increases memory latency, but also becomes a primary bottleneck in multiprocessor configurations as processors compete for the shared bus. In fact, studies have shown that Intel's Xeon and Itanium bus architectures can even be saturated by a single processor alone². As processors are added to this

1. Johan De Gelas, "Socket 939: New Athlon 64s", Aces Hardware, 2004, <http://www.aceshardware.com/read.jsp?id=65000306>
2. Cornell Theory Center, "Benchmarking Intel Systems and Understanding Results", <http://www.microsoft.com/windows2000/docs/hpcintbench.doc>

shared bus architecture, the available bandwidth is quickly approached, consequently resulting in minimal performance gains.

4.0 The UltraSPARC T1 Processor as a Server-on-a-Chip

What has traditionally been implemented in SMP systems is now implemented on a single chip.

The UltraSPARC T1 processor's high memory bandwidth is simply a consequence of the CMT design of building a server-on-a-chip. The use of an SMP servers composed of multiple processors designed to maximize performance of commercial server workloads has been a very expensive proposition -- both the cost and power to operate SMP servers are at all-time highs. A more efficient approach is to build a machine using simple cores aggregated on a single die, with a shared on-chip cache and high bandwidth to large off-chip memory, thereby aggregating an SMP server on a chip. This has the added benefit of low-latency communication between the cores for efficient data sharing in commercial server applications.

Consider the difference in system bandwidth between a traditional SMP server and an UltraSPARC T1 processor shown in Figure 2.

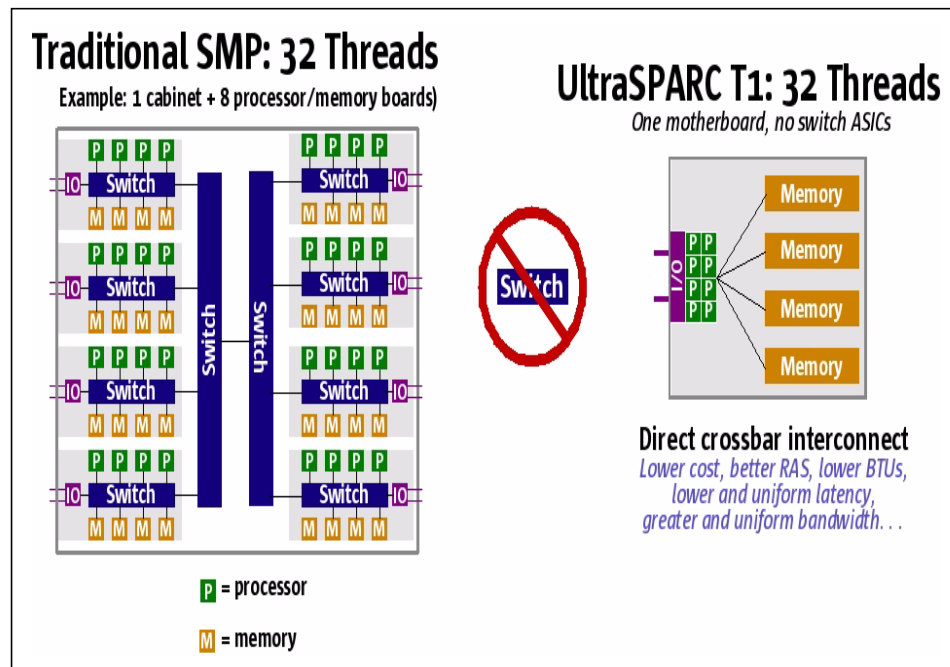


FIGURE 2 Traditional SMP Server vs. UltraSPARC T1 Processor

In conventional SMP servers using multiple single core processors with coherent system interconnects, coherence misses go out over low-frequency off chip buses and links, and can have high latencies. The UltraSPARC T1 processor design with its many multithreaded cores, shared on-chip cache, and on-chip crossbar interconnect eliminates these misses and replaces them with low latency communication. With the UltraSPARC T1 processor's high bandwidth design and high levels of integration, what has been traditionally implemented in SMP systems is now possible using a single chip. The intricacies of multiple processors communicating over an interconnect in an SMP system is now implemented by multiple cores communicating over an interconnect -- all on a single chip -- at much closer proximity.

The radical difference in system bandwidth between an SMP server and a server-on-a-chip like the UltraSPARC T1 processor is compared in Table 4.

TABLE 4 Relative Bandwidth Per Core: Enterprise 10K vs. UltraSPARC T1 Processor

	Total Number of Cores	L2 Cache to Main Memory Bandwidth (in Gbytes/s)	Bandwidth Per Core (in Gbytes/s)	Relative Bandwidth Per Core
Enterprise 10K	64	12.8	0.2	1.0X
UltraSPARC T1 Processor	8	76.8	9.625	>48.0X

For the same performance, the UltraSPARC T1 processor delivers 8 times less the number of cores and 2 times less the number of threads than a 64-processor E10K.

This comparison focuses on the event when a processor does not have the data item it needs in the L2 cache and must access the next level of the memory hierarchy (namely main memory in this example). The Enterprise 10K (E10K) interconnect supports bandwidths of up to 12.8 Gbytes/s. Configured with 64 processors (or 64 cores), the E10K allocates 200 Mbytes/s of bandwidth per processor. In sharp contrast, the UltraSPARC T1 processor's L2 cache to main memory bandwidth is 76.8 Gbytes/s, supporting a staggering 9.625 Gbytes/s per core, or almost 50 times more bandwidth per core than the E10K¹. Provided with this extremely high bandwidth interconnect, the UltraSPARC T1 processor was able to deliver the same performance as the E10K, with 8 times less the number of cores and 2 times less the number of threads. Although the difference in interconnect bandwidth was only considered in this example,

1. A single 400 MHz UltraSPARC II is equivalent to an UltraSPARC T1 processor thread. For OLTP applications, using actual test results for the E10K and performance estimates for the UltraSPARC T1 processor, a single the UltraSPARC T1 processor-based system is capable of delivering similar performance to the E10K.

the UltraSPARC T1 processor also offers orders of magnitudes in savings in cost, power dissipation, real estate, and component count.

5.0 Conclusion

The advent of CMT processors is first realized by the UltraSPARC T1 processor. Designed from the ground-up, the UltraSPARC T1 processor strives to achieve maximum IPC with a balanced system comprised of a crossbar interconnect, caches, and a memory interface. Combined with its high bandwidth and high concurrence of executing many instructions simultaneously, the UltraSPARC T1 processor is the industry's most innovative and effective solution.